

## **HUMAN-MACHINE TEAMING FOR DYNAMIC FAULT MANAGEMENT IN NEXT-GENERATION LAUNCH VEHICLES<sup>1</sup>**

R. S. McCann  
NASA Ames Research Center  
Moffett Field, California

&  
J. McCandless  
San Jose State University Foundation  
Moffett Field, California

### **ABSTRACT**

In the thirty years since the shuttles started flying, vehicle health management technologies have advanced to the point where they can autonomously monitor systems functioning and detect, diagnose and respond to systems malfunctions in real time. On a crewed vehicle, however, we argue that it is neither advisable nor practical to fully automate the health management system. Instead, human and machine intelligence should function as an integrated team, with roles and responsibilities carefully allocated to optimize the functioning of the human/machine system. In this paper, we first describe how dynamic fault management occurs in the shuttle today. We then develop a "straw man" operational concept for fault management in a hypothetical future version of the shuttle equipped with autonomous fault management technologies. The concept includes a description of human-machine functional allocation and associated human-computer interfaces. We illustrate the operational concept with an example of cooperative management of a disturbance in one of the main engine's helium supply system.

### **INTRODUCTION**

Operating a flight vehicle is among of the most challenging of all human occupations. The classic "aviate, navigate, communicate" dictum captures one of the most important sources of difficulty, namely the multi-tasking nature of the environment. Cockpit tasks fall into two categories. Discrete tasks are triggered by a particular event and have a fixed duration. An example is a radio request from air traffic control for a course change, which the pilot acknowledges and carries out by reprogramming the flight management computer. Another example, dynamic fault management<sup>1</sup> is the focus of this paper. A sudden perceptual event (typically an alarm) alerts the crew to the presence of a disturbance in the functioning of an onboard system. The crew must diagnose the source of the disturbance and take appropriate mitigating actions. The second category encompasses tasks of a more continuous nature, such as activities needed to maintain a high level of situation awareness of the vehicle's navigation state. These activities typically take the form of a sequential acquisition (scan) of information from flight instruments, cockpit displays, and the out-the-window view. Both laboratory studies<sup>2</sup> and analyses of crew performance in the cockpit<sup>3</sup> have identified multi-tasking as a major source of workload and human error. Fortunately, as the aircraft industry has matured, multi-tasking demands on the crew have generally decreased. There are two factors contributing to the reduction. First, aircraft designers have taken full advantage of modern computing power to automate traditional cockpit tasks, such as active flight control. Second, aeronautical engineers have been developing, testing, and refining aircraft systems for over a century. By now, these systems are so reliable that dynamic fault management is a very rare requirement.

---

<sup>1</sup> Approved for public release, distribution is unlimited.

This work was supported by NASA's Space Human Factors Engineering Program (RTOP # 131-20-30).

Unfortunately, the same level of systems reliability does not exist on the shuttles. Particularly during the dynamic phases of flight (ascent and entry), crew safety and mission success are critically dependent on the nominal functioning of a large number of shuttle systems. Several of these (most notably the propulsion system) are far more complex and physically dynamic than their aircraft counterparts, and operate in a much harsher physical environment. Combined with the fact that the shuttle is still an experimental flight vehicle, with a relatively short history of development, it is no surprise that systems malfunctions are far more likely to occur (and when they do occur, are potentially far more dangerous) than on a modern aircraft.

Recognizing this fact, the astronaut-training program includes extensive instruction in the nature and functioning of shuttle systems. Working in ground-based simulators, astronaut candidates also receive extensive training in dynamic fault management so that they become adept at recognizing fault signatures and working the associated procedures. By the time they fly, crewmembers are almost as knowledgeable about shuttle systems functioning as the subject matter experts at mission control.

Even with the training, however, dynamic fault management is a time-consuming and difficult activity.<sup>4</sup> On-board sensors record subsystem parameters such as temperature, pressure, speed, flow rate, and system configuration. These values are displayed directly to the crew on dedicated meters and electronic system summary displays. Sensor data are also routed through limit-sensing software whose function is to alert the crew to the presence of an out-of-limit value. The alert typically takes the form of an auditory alarm and written fault message. Unfortunately, the caution and warning software cannot distinguish between real malfunctions and false alarms generated by failed sensors. Furthermore, shuttle systems are highly interconnected and interdependent, such that the effects of a malfunction in one component can propagate rapidly into other components, producing a cascade of alarms and messages that make it difficult to discern the "root cause" of the problem. For example, if a power distribution assembly fails, all physical devices powered by the assembly, such as fans and pumps, will fail and generate their own alarms and fault messages.

These problems force the crew to engage in time consuming information-gathering activities to identify false alarms and/or make root cause determinations. These activities can include cross-checking parameter values on systems summary displays, resetting parameter values to determine if the malfunction re-occurs in a timely manner (a check for sensor failure), checking for related systems malfunctions, and performing inventories of failed equipment. Once the malfunction is confirmed, a crewmember has to locate the appropriate set of fault management procedures on cue cards or in paper versions of the vehicle flight data files. The procedures are written in the form of a terse set of abbreviations, acronyms, symbols and instructions that require considerable training to decipher. Then, after deciphering the code and determining the appropriate actions, the crewmember must locate the correct switch(es) or dial(s) on a plethora of control panels located to the left, overhead, and to the right of the crew station. Finally, once the appropriate actions are performed, indicators must be monitored to ensure that the reconfiguration has had the desired result.

Laboratory studies of fault management behavior<sup>5,6</sup> have found that controllers of complex engineering systems often do not (or are unable to) time-share fault management activities between two or more overlapping malfunctions. Instead, the controllers work the malfunctions strictly sequentially. There is also evidence<sup>7</sup> that this phenomenon, known as cognitive lockup<sup>8</sup>, can occur to flight crews when faced with a systems malfunction, particularly when the malfunction occurs in an unusually stressful or high-workload context. In one of the clearest examples, from the National Transportation Safety Board archives, cognitive lockup was manifest as a failure to time-share dynamic fault management activities with the more continuous (e.g., instrument scanning) tasks to maintain crewmembers' awareness of the vehicle's navigation state. The aircraft eventually crashed into the Florida Everglades. On a shuttle mission, the vehicle is operating so close to performance limits during ascent and entry that propulsion and navigation states must be monitored extremely carefully. The possibility of cognitive tunneling on a difficult fault management problem represents a clear and present hazard to crew safety and mission

success. Therefore, there are compelling reasons to both simplify and shorten the duration of dynamic fault management processes.

NASA's Johnson Space Center has recently completed a redesign of the shuttle cockpit displays and caution and warning system to make the cockpit more user friendly. The upgrade includes several features to reduce workload associated with dynamic fault management activities.<sup>9</sup> Most notably, a comprehensive rule-based "expert system" has been developed to provide an automated root-cause fault determination capability. When implemented (several years from now), the system will function as a filter between limit-sensing software and the cockpit, allowing only those annunciations associated with the root cause itself to annunciate and generate a fault message. Indications associated with the "children" of the root cause will be inhibited. The filter also incorporates phase-of-flight knowledge to inhibit nuisance alarms. Together with a redesigned fault summary display, this enhanced caution and warning system should make fault diagnosis considerably less difficult and less time consuming than it is today.

Another focus of the cockpit redesign has been the system summary displays. Figure 1 illustrates the major features of the display redesign with reference to the current and redesigned displays showing the health and status of the main engines. The most obvious format change is the replacement of columns and rows of digital sensor values with pairings of digital values and graphical elements. These elements are spatially arranged to give crewmembers "at a glance" information concerning system configuration (e.g., valve status), operation (e.g., flow) and the functional interrelationships between subcomponents<sup>10</sup>. As we shall see shortly, displays that graphically depict subsystem components are crucial to support effective human-machine teaming for dynamic fault management.

In the existing shuttles, automated fault management capabilities are largely limited to the annunciation of out-of-limit sensor values. In the upgraded cockpit, the primary augmentation to existing automation will be the root-cause fault determination capability of the enhanced caution and warning system (though it is important to note that other tasks, such as dynamic flight management, will also benefit from enhanced automation). In the broader aerospace industry, vehicle health management systems have been developed with much more extensive capabilities. For example, the MD-11 boasts systems controllers that manage certain classes of faults completely autonomously; the crew is simply notified as to what the fault was, and what actions were taken by the controller. In 1999, NASA launched Deep Space 1, an experimental unmanned spacecraft whose charter was to demonstrate new technologies. Although most of NASA's attention was focussed on the innovative propulsion system, Deep Space 1 also successfully flight-tested a comprehensive health management system that could both diagnose malfunctions, by making inferences based on systems-level models, and work the associated procedures in a fully autonomous fashion. The technology used on Deep Space 1 is only one of a number of advanced model-based diagnostics and reasoning health-management tools that have been developed and flight tested recently.<sup>11</sup>

In next-generation-manned spacecraft, stringent requirements for crew safety and operational efficiency are virtually mandating the incorporation of these advanced health management technologies. From an operational perspective, should the goal then be to design the human out of the health management loop entirely, giving full responsibility to machines? There are several compelling reasons why this is not appropriate. Until software validation and verification techniques become more robust, entrusting health management of a vehicle as complex as a spacecraft entirely to a software system is inherently risky. The risk is exacerbated when the software resides on space-based platforms, as these are more vulnerable than ground-based platforms to radiation-induced "single event upsets." Beyond the obvious issue of machine reliability, however, when astronauts are onboard the vehicle, a fully automated design solution is simply suboptimal. As we noted earlier, when it comes to spacecraft systems, crewmembers are subject matter experts. Taking them out of the loop entirely amounts to a decision to waste valuable onboard expertise. Equally important, humans and machines bring different capabilities and different vulnerabilities to bear on the fault management process. A full discussion of this topic is beyond the scope of the current article; the interested reader is referred to an earlier citation (<sup>7</sup>) for

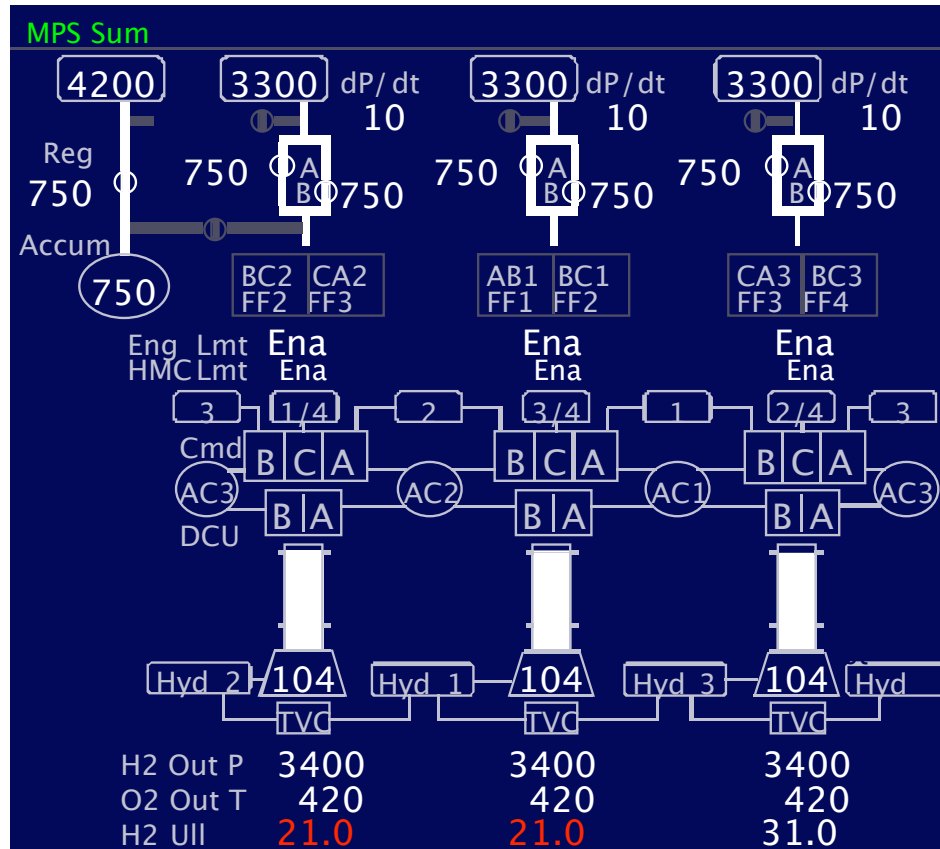
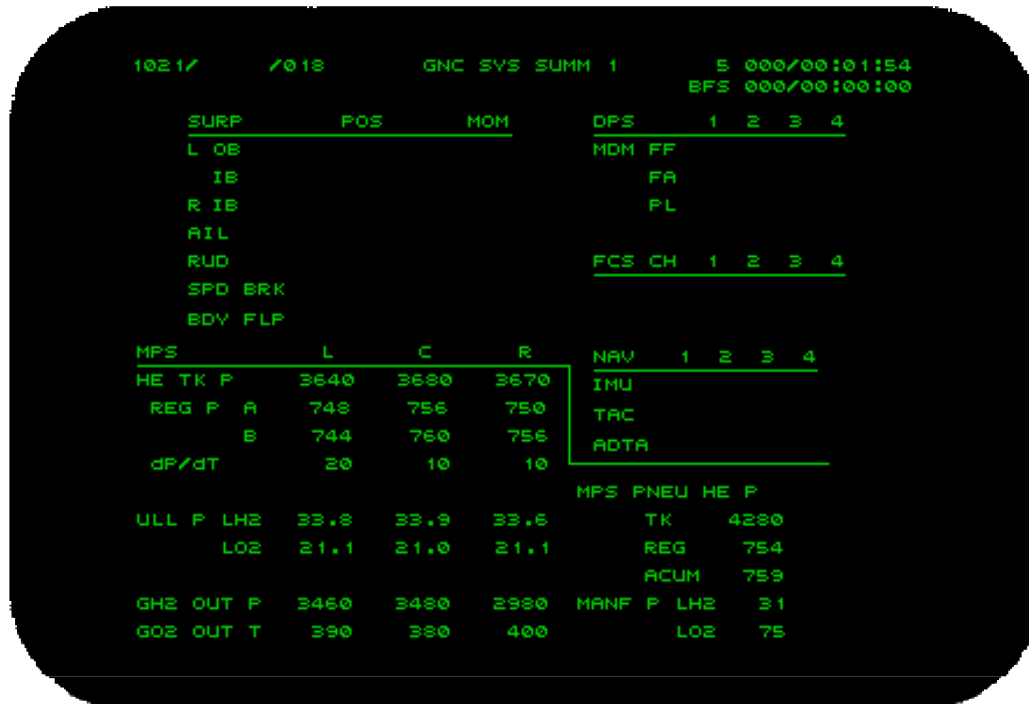


Figure 1. Example of the shuttle cockpit upgrade display format redesign. Upper panel: BFS GNC SYS SUM 1 display. Lower panel: Proposed main propulsion system summary display.

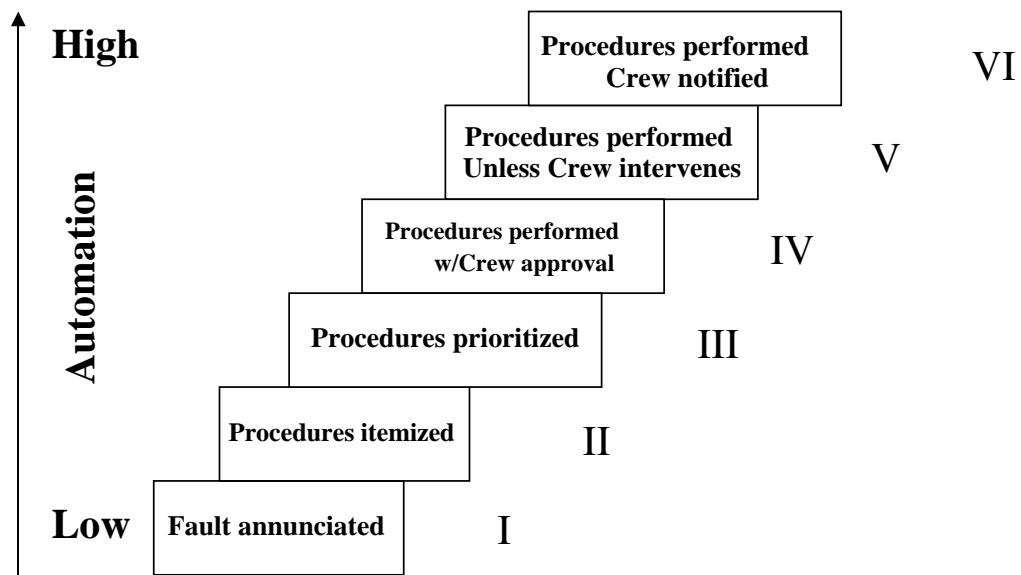
extensive discussion. We do want to point out that human and machine capabilities are frequently complementary, with strengths in one offsetting weakness in the other. For example, humans have a fluid reasoning capability that makes them effective problem-solvers in novel situations; by contrast, automated systems are exceedingly brittle in the face of the unexpected. On the other hand, humans function in a nondeterministic fashion; if a human is faced with a set of procedures to work, various memory and attentional factors can lead him or her to omit procedures, or perform them inaccurately. Automated systems are deterministic: if the software is accurately coded, automation will consistently perform a set of procedures, in the correct order, without error.

From a variety of perspectives, then, the optimal design solution for a next-generation crewed space vehicle is for people and intelligent systems to "combine forces" and work fault management activities in a collaborative fashion. Fortunately, the major design issues that need to be addressed for effective human-machine "teaming" are now well documented.<sup>12</sup> These issues include:

- Ensuring visibility into automated functioning. Automation is deemed "clumsy"<sup>13</sup> if the workings of the automation are opaque to the human. In the case of dynamic fault management, a crewmember might be unsure of the basis for an automated fault diagnosis, why the machine is taking the actions it is, or the consequences of machine actions for systems mode and systems functioning.
- Providing a "meta-level" monitoring capability. It is crucial that the crew be able to monitor the health of the health management system itself. That way, if the management system malfunctions, the crew is quickly able to recognize this fact, and take over health management responsibilities themselves.
- Minimizing the "out-of-the loop unfamiliarity" (OOTLUF) problem while still ensuring sufficient "value added" by the automation. The functional allocation between human and automation needs to strike a balance between these two considerations. Humans should be involved in a way that promotes an adequate level of situation awareness (SA) of the problem and the actions being taken in response to it. If too much of the process is automated, humans can lose SA, with serious consequences.<sup>4,14</sup> On the other hand, enough functions have to be automated that the automated system "buys its way" onto the spacecraft.

These considerations force us once again to confront the issue of what machines do better than humans, and vice versa. In the case of dynamic fault management on a manned spacecraft, we believe that the appropriate targets for automation are fault diagnosis and retrieving and working the appropriate procedures.

Accordingly, we offer the following "straw man" operational concept for fault management in a hypothetical upgraded version of the shuttle equipped with an integrated advanced health management system. We assume that the health management system has the ability to autonomously monitor system state, diagnose malfunctions, and work malfunction procedures in real time. Figure 2 (from McCann & McCandless, 2002; citation 4) modifies the classic Sheridan-Verplank (S-V) automation scale<sup>15</sup> to fit the dynamic fault management realm. Level 1 corresponds to the current level of automation in the shuttle, where parameter values that go out of limits are annunciated automatically. With the addition of root-cause determination (cockpit upgrade), "fault annunciated" also includes a measure of automated diagnosis capability. We assume that diagnostics capabilities would be further enhanced with the advanced health management tool set. Levels 2 and 3 incorporate automated retrieval and presentation of the appropriate procedures (an electronic flight data file), such as those found today on Boeing 747-400's. Level IV and above automate the individual procedures. In our operations concept, Level IV is the "default" or "nominal" case; the automation diagnoses the fault and performs each procedure upon receipt of a "permission" from a crewmember. This functional division frees the crew from having to locate procedures in paper flight data files, and from having to physically perform the procedures. However, by requiring sequential consent for each procedure, the human is "on the hook" to verify the



**Figure 2. Levels of automation (LOA) for dynamic fault management<sup>4</sup>**

automated diagnosis, and retains ultimate control over the fault management operation. These responsibilities ought to ensure a high level of situation awareness. The sequential "step-through" nature of the procedures working should also help to maintain a "shared mental model" between human and automation of the problem, the reconfiguration actions that are being taken in response to it, and the response of the system to these actions.

Having selected a candidate functional allocation, the remaining challenge is to design the supporting human-computer interfaces. Our departure point for this discussion is a malfunction in the helium supply system for the space shuttle's left engine, described below. We will start with a description of how the crew would handle this malfunction in the upgraded shuttle cockpit, using the main propulsion system (MPS) display (Figure 1, bottom panel) as the primary point of reference. We then describe the display modifications needed to provide effective crew interaction with S-V level IV automation.

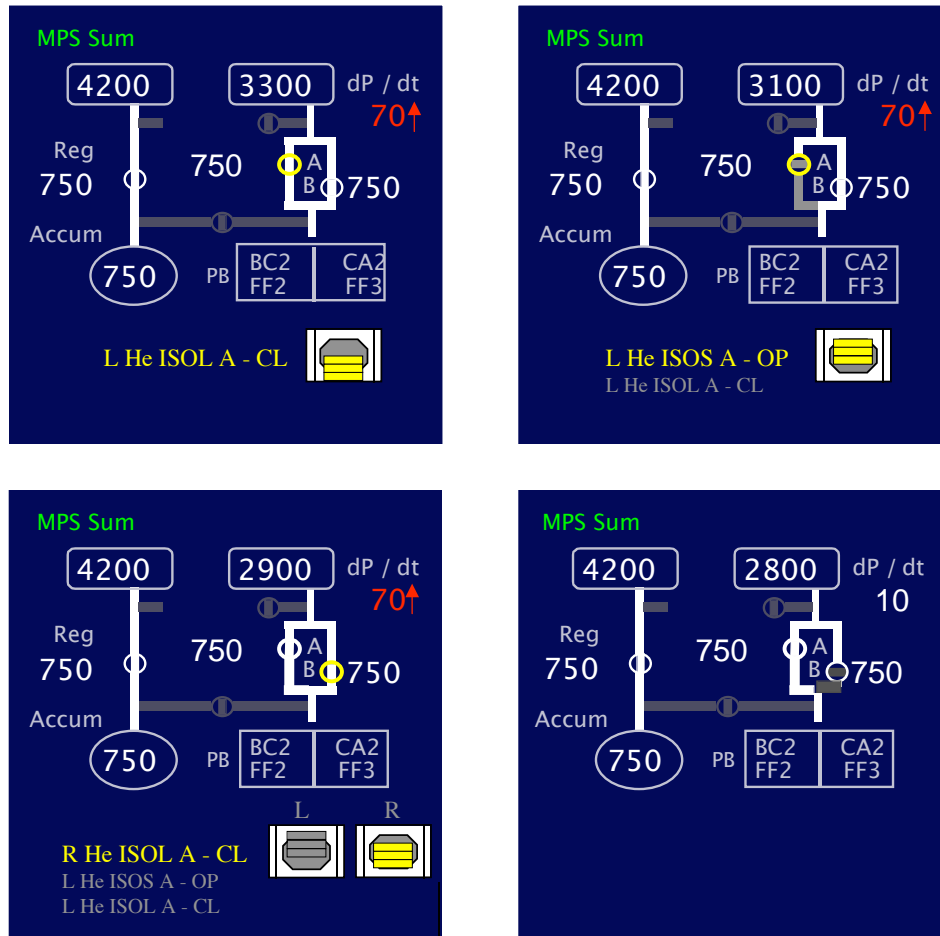
Briefly, each SSME requires a continuous supply of helium to pressurize an intermediate seal in the high-pressure oxidizer turbopump. Helium is supplied from tanks that are collectively depicted on the MPS summary display (Figure 1; bottom panel) as the topmost rectangles with the digital tank pressure readings contained inside. Helium flow, indicated by the  $dP/dt$  (change in tank pressure over time) digital value to the right of the tank, proceeds from the tanks through two redundant legs, A and B, each with a regulator that reduces pressure to 750 lbs. per square inch. The regulated pressure in each leg is also depicted with digital values. From there, the two legs rejoin to form one supply line, and regulated helium flows through the supply line into the left engine. Note also the depiction, to the left of the left engine supply system, of an additional helium supply system. This "pneumatic helium supply" system can provide additional helium to supplement the supply to any main engine in the case of a supply system leak, or to shut down the main engines via application of gaseous (pneumatic) pressure in the event of a hydraulic failure. An interconnect line exists downstream of the regulators so that regulated helium from the left engine supply system can feed the pneumatic system, when the need arises.

Various aspects of the display provide useful representational aiding.<sup>10</sup> Supply lines can take on one of two brightness levels, either gray (signifying no flow) or bright white (indicating flow). Valves are depicted by the small circles superimposed on the supply lines, together with the section of the line

embedded in the circle. If the valve is closed, the embedded section is drawn perpendicular to the line itself, "breaking" the line perceptually, and both the embedded line section and the embedded portion are gray (no flow). If the valve is open, the embedded section is aligned with the rest, completing it perceptually, and is rendered bright white. Off-nominal parameter readings are indicated with color-coding. Nominal values are gray or white, off-nominal values are yellow or red, and missing values are cyan.

Consider a malfunction in the left engine helium supply system that results in the rate of flow (change in pressure over time, or  $dP/dt$ ) going off-nominal high. The indications to the crew might include an auditory alarm and a fault message of the form "MPS He P" at the bottom of the MPS display. In addition, the  $dP/dt$  digital value just to the right of the left engine tank graphic would be colored red and accompanied by a red "up" arrow. This set of symptoms is associated with several possible malfunctions. There could be a leak in the tanks themselves, a leak in Leg A or Leg B, or a Leg A or Leg B regulator failure. Suppose the actual problem is a leak in Leg B. Having read the fault message, the pilot would first locate the fault management procedures on either a cue card, or in the appropriate paper flight data file. The procedures guide the pilot through an abductive reasoning exercise<sup>16</sup> to isolate and fix the problem. The first step is to close the Leg A isolation valve by altering the position of a switch on a panel to the right of the pilot's right leg. If, in response to this action,  $dP/dt$  returned to a nominal value, the action would have successfully associated the problem with either a leak in the A leg or an A regulator failure; either way, the system would have been "safed". Since this is not the case in our example,  $dP/dt$  would not lessen as a consequence of this action, and would stay red. The next procedures are to re-open the Leg A isolation valve, and then close the Leg B isolation valve. In our hypothetical example, these actions would isolate the leak, and would solve the high flow problem. The crew would note the return of  $dP/dt$  to a nominal value (the digital values would turn back to white), and the pilot would exit the procedures (if, however, closing Leg B had not solved the high flow problem, the source would have been a "nonisolatable" leak, and the left engine would have had to be shut down).

How might this process work with Level IV automation? The display concept is illustrated in Figure 3. The idea is that the MPS display morphs into the display shown in the top right panel of Figure 3, which selects, highlights, and magnifies the section of the system where the malfunction has occurred. Again, the left engine flow indicator ( $dP/dt$ ) is colored red and has an up arrow beside it, indicating the off-nominal high flow condition. The automation retrieves the appropriate set of procedures and presents a text version of the initial procedure at the bottom of the display, commanding the Leg A isolation valve to be taken to the closed (down) position. A switch-like graphic is displayed beside the text, with the switch position indicator showing the switch position commanded by the text message. Graphically, the "Close Isolation Valve A" command is also indicated by coloring the circle<sup>17</sup> surrounding the valve segment yellow. This graphical indication is linked perceptually to the text command and the actual action by also coloring the text and the switch position on the virtual switch icon yellow. The pilot gives permission for the automation to perform the switch throw by physically touching the yellow switch position indicator on the display (although a verbal permission should also be considered, assuming a sufficiently reliable voice recognition system). Once the action has been carried out, the display would shift to the configuration in the upper right panel. The graphics provide immediate feedback as to new system status; the Leg A isolation valve is now closed, and the flow through Leg A has been halted (gray color). However,  $dP/dt$  is still red, indicating that this action has not solved the high flow problem. Therefore, the first procedure in the textual procedures list is now de-emphasized (gray), and the next procedure in the sequence appears at the top of the list (again, colored yellow). This procedure calls for the Leg A valve to be opened back up. Graphically, the orientation of the line segment inside the circle shows current valve position (closed), while the yellow color of the circle indicates that a change in valve status (to open) is being commanded. The pilot again gives permission for this action by touching the yellow switch position indicator, and the display shifts to the configuration in the bottom left panel. As is clear from the panel, the Leg A isolation valve is now open, there is helium flow through Leg A, and the automation has retrieved the next command from the electronic flight data file, which is to close the Leg B (right leg) isolation valve. Again, graphically, in addition to the yellow text line and the yellow color of the switch position indicator, coloring the circle component of the Leg B valve indicator yellow indicates a



**Figure 3. Illustration of human-computer interface for cooperative dynamic fault management. Upper left panel: Left (Leg A) isolation valve commanded "Closed". Upper right panel: Left (Leg A) isolation valve commanded "Open". Lower left panel: Right (Leg B) isolation valve commanded "Closed". Lower right panel: Left Engine Helium Supply System configuration upon exiting procedures.**

commanded change in valve position. The crewmember, after due consideration, gives permission for this action to be performed, and the display morphs once again into the configuration in the lower right panel. Here, finally, dP/dt has dropped back to a nominal value, and is no longer colored red. This signals that the current system configuration (Leg A isolation valve Open, Leg B isolation valve Closed) has isolated the problem and "safed" the system. The final configuration is available at a glance from the graphics.

There are a number of features to this human-computer interface that are worth pointing out. First, the close coupling between the graphics and the text messages allows the crew to cross check the command action with the current system configuration. This kind of cross checking can be valuable for preventing errors. For example, helium flow into the turbopump must be continuous; any interruption could quickly bring the gaseous fuel/oxidizer mixture in contact with the liquid oxygen, with explosive results. Thus, the one configuration that must be avoided is having both Leg A and B valves closed at the same time. In the sequential approach, illustrated in Figure 3, closing Leg B is not even provided as an option until Leg A has been opened back up. The graphical cues showing Leg A valve status open, and



the presence of flow through Leg A, provide a further confirmation that it is safe to proceed with the closing of isolation valve B.

Of course, an operations concept is only that; a concept. In order for the concept to be accepted by the mission operations and astronaut communities, it must be extensively tested and validated. It must also demonstrate a sufficient improvement over current operations to make the underlying development costs worthwhile. We are currently developing a part-task shuttle cockpit simulator at NASA Ames Research Center to evaluate these concepts. Our approach is to simulate dynamic phases of shuttle flight and to interject carefully selected systems malfunctions that the operator must diagnose and work in real time. The operator's eye movements are recorded at a rate of 60 Hz. We combine the eye movement data with a recording of operator actions, such as switch throws and display navigation activities, to build a comprehensive "picture" of human-machine fault management activity. Quantifying the temporal duration of individual components of that activity in today's shuttle cockpit environment provides baseline measures of fault management performance. In future, these measures will then be used to quantify improvements (such as reductions in total fault management duration, and degree of time sharing of fault management activities with concurrent cockpit tasks) associated with operational concepts such as the one described in this paper.

Our operational concept has focussed on the role of an advanced health management system in diagnosing malfunctions and working with humans on the procedures. However, incorporating a full suite of advanced health management technologies into next generation spacecraft may require a more extensive redesign to the system summary displays than described here. Such a suite would almost certainly include advanced data-processing algorithms that assess (categorize) the functioning of a system and its components by performing real-time trend analyses on sensed data. These tools have interesting implications for the design of systems summary displays. For example, it may be appropriate to replace real-time digital readouts of parameter values with some form of "inference indicator" that communicates the algorithm's current "belief" about the health of the system's components.

## REFERENCES

1. Woods, D. (1994). Cognitive Demands and Activities in dynamic fault management: Abductive reasoning and disturbance management. In N. Stanton (Ed.), *Human Factors in Alarm Design* (pp. 63-92). London: Taylor and Francis.
2. Rogers, R. D., & Monsell, S. (1995). The cost of a predictable switch between simple cognitive tasks. *Journal of Experimental Psychology: General*, 124, 207-231.
3. Chou, C. D., Madhavan, D., & Funk, K. (1996). Studies of cockpit task management errors. *The International Journal of Aviation Psychology*, 6, 307-320.
4. McCann, R. S., & McCandless, J. (2002). Toward a crew-system concept for real-time fault management in next-generation aerospace vehicles. In *Proceedings of the Joint Army-Navy-NASA-Air Force 2nd Modeling and Simulation Subcommittee Meeting*, Sandestin, FL.
5. Moray, N., & Rotenberg, I. (1989). Fault management in process control: Eye movements and action. *Ergonomics*, 32, 1319-1342.
6. Moray, N., Inagaki, T., & Itoh, M. (2000). Adaptive automation, trust, and self-confidence in fault management of time-critical tasks. *Journal of Experimental Psychology: Applied*, 6, 44-58.

7. Rogers, W. H., Schutte, P. C., & Latorella, K. A. (1996). Fault management in aviation systems. In R. Parasuraman & M. Mouloua (Eds.), *Automation and Human Performance: Theory and application* (pp. 281-317). New Jersey: Erlbaum.
8. Sheridan, T. B. (1981). Understanding human error and aiding diagnostic behavior in nuclear power plants. In J. Rasmussen & W. B. Rouse (Eds.), *Human detection and diagnosis of system failures* (pp. 19-36). New York: Plenum Press.
9. McCandless, J., McCann, R. S., & Hilty, B. (2003, in press). Upgrades to the caution and warning system of the space shuttle. In *Proceedings of the Human Factors and Ergonomics Society 47th Annual Meeting*.
10. Guerlain, S (2002). The MPC Elucidator: A case study in the design for human-automation interaction. *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans*, 32, 25-40.
11. Patterson-Hine, A., Hindson, W., Sanderfer, D., Deb, S., & Domagala, C. (2001). *A model-based health monitoring and diagnostic system for the UH-60 Helicopter*. Presented at the American Helicopter Society 57th annual forum. Washington, DC.
12. Malin, J., & Schreckenghost, D. (1992). Making intelligent systems team players: Overview for Designers. *NASA Technical Memorandum 104751*.
13. Billings, C. E. (1997). *Aviation automation: The search for a human-centered approach*. Hillsdale, NJ: Erlbaum.
14. Endsley, M. R., & Kiris, E. O. (1995). The out-of-the loop performance problem and level of control in automation. *Human Factors*, 37, 381-394.
15. Sheridan, T. B. (1992). *Telerobotics, Automation, and Human Supervisory Control*. Cambridge, MA: The MIT Press.
16. Woods, D. (1995). The alarm problem and directed attention in dynamic fault management. *Ergonomics*, 38, 2371-2393.
17. Malin, J., Kowing, J., Schreckenghost, D., Bonasso, P., Nieten, J., Graham, J. S., Fleming, L. D., MacMahon, M., & Thronesbery, C. (2000). Multi-agent diagnosis and control of an air revitalization system for life support in space. In *Proceedings of the IEEE Aerospace Conference*.